arXiv:1704.02004v2 [q-bio.PE] 17 Apr 2017

TOPOLOGICAL METRIZATIONS OF TREES, AND NEW QUARTET METHODS OF TREE INFERENCE

JOHN A. RHODES

Abstract. Topological phylogenetic trees can be assigned edge weights in several natural ways, highlighting different aspects of the tree. Here the rooted triple and quartet metrizations are introduced, and applied to formulate novel fast methods of inferring large trees from rooted triple and quartet data. These methods can be applied in new statistically consistent procedures for inference of a species tree from gene trees under the multispecies coalescent model.

1. Introduction

By the oxymoronic term *topological metrization* of a tree, we mean an assignment of edge weights to the tree that reflects solely topological information about the tree. Such a metrization adds no new information to the tree, and at first might appear to be merely a curiosity. Surprisingly, however, two such topological metrizations play key roles in fast, statistically consistent methods of species tree inference in phylogenomics. These metrizations, as will be explained, can be viewed as emphasizing the clades or splits displayed on a rooted or unrooted tree.

In this short note, two new topological metrizations are introduced, built upon a tree's displayed rooted triples (induced 3-leaf rooted trees) and displayed quartets (induced 4-leaf trees) for rooted and unrooted topological trees, respectively. These metrizations have an important and useful feature: Intertaxon distances can be computed from a list of rooted triples or quartets displayed on the tree — without knowing the tree itself. Even if the quartets believed to be on the tree have some error, approximations to the intertaxon distances on the tree are still obtained. Thus the tree can be quickly inferred by any of the well-known distance-based methods for tree building or selection that are robust to error. This approach extends to give new statistically consistent methods of species tree inference from samples of topological gene trees drawn from the multispecies coalescent model, adding to those few already known [DDBR09, ADR13, ADR17].

The simplest topological metrization of an unrooted tree assigns weight 1 to all edges of the tree, so intertaxon distance along the tree is the standard graph-theoretic distance between leaves. Though this has often been used for heuristic purposes, a *non*-heuristic application arises in the NJ_{st} method of species tree inference from gene trees [LY11], which has been reimplemented in the ASTRID software [VW15], and further studied and renamed the U-STAR (Unrooted Species Tree inference by Average Remetrization) method in [ADR17]. This method takes a sample of unrooted topological gene trees,

Date: April 19, 2017.

metrizes them with unit edge weights, encodes the metric trees through pairwise inter-taxon distances, and computes the arithmetic mean distances over the gene trees (i.e., applies average consensus [LC97, Bry03] to the metrized gene trees). For enough data produced from the multispecies coalescent model, this average distance closely fits some metrization of the species tree, so statistically consistent inference of the unrooted species tree topology is achieved by a tree-building algorithm such as Neighbor Joining (NJ), or other methods [ADR17].

For species tree inference from rooted topological gene trees, the STAR method [LYPE09] only differs from U-STAR by using a slightly different metrization of the gene trees. Since STAR and U-STAR inference have been shown to depend not on the full gene tree sample, but only on the clade and split frequencies in the sample [ADR13, ADR17], the STAR and U-STAR metrizations of topological gene trees can be viewed as natural ones that highlight certain topological features — clades, or splits — of the trees through the assigned edge weights.

These ideas motivate the exploration here of other ways in which phylogenetic trees might be metrized to reflect only their topological features, and how such metrizations might be used for inference. Basic definitions are given in Section 2. The clade and split metrizations are presented in detail in Section 3. The primary theoretical contributions of this note are the two new metrizations, associated to rooted triples and quartets, that are developed in Sections 4 and 5. The applicability of this theory to supertree inference from rooted triple and quartets and to species tree inference from gene trees is then developed in Sections 6 and 7. Section 8 concludes with a few general comments.

While very limited simulations in Sections 6 and 7 illustrate the potential of these methods to address empirical problems, the emphasis throughout later sections is on presenting the inference framework rather than exploring its performance. Extensive testing through well-designed simulation studies is still needed.

2. Preliminaries

Throughout this work, we let X denote a finite set of N taxa. We use upper case letters $A, B, \ldots$ for subsets of X, and lower case letters $a, b, \ldots$ for elements of X.

A *split of* X is a bipartition $X = A \sqcup B$ of the taxa into non-empty subsets, and is denoted $A|B = B|A$. A *clade* of X is a non-empty subset $A \subseteq X$ of the taxa. A *rooted triple of* X is a subset of three elements of X, partitioned into a pair a, b and a singleton c, and denoted $ab|c = ba|c$. A *quartet* of X is a subset of four elements of X, partitioned into two pairs a, b and c, d, and denoted $ab|cd = ba|cd = \cdots = cd|ab$.

Suppose the taxa X bijectively label the leaves of a rooted tree T^r, or of an unrooted tree T, with the root of degree at least 2 and all other internal nodes of degree at least 3. Then T^r and T are said to be *phylogenetic trees on* X. All edges on a rooted tree T^r are directed away from the root, so, for instance, the root is ancestral to all leaves. Edges on an unrooted tree T are undirected. A phylogenetic tree is *binary* if the minimal degree conditions on the nodes are met.

We say T^r *displays the clade* A if the most recent common ancestor (MRCA) on T^r of the taxa in A has as its descendants in X precisely the set A. Thus clades displayed on

a rooted tree correspond to its nodes, and if the tree is binary, it displays exactly $2N-1$ clades, including all singleton clades and the clade X. We say T^r *displays the rooted triple* $ab|c$ if the MRCA of a and b is a proper descendent of that of a, b, and c. If $N \geq 3$, a binary rooted tree on X displays $\binom{N}{3}$ rooted triples, one for each choice of three taxa.

Similarly, for an unrooted tree T on taxa X, we say T *displays the split* $A|B$ if the bipartition is obtained by removing an edge of T and partitioning X according to connected components of the resulting graph. If $N \geq 2$, a binary tree displays $2N-3$ splits. We say T *displays the quartet* $ab|cd$ if on the induced 4-leaf tree relating a, b, c, and d there is an internal edge separating the taxa a and b from c and d. If $N \geq 4$ a binary tree thus displays $\binom{N}{4}$ quartets, one for each subset of four taxa.

Suppose positive weights are somehow assigned to the edges of T^r or T, so the tree is now a *metric tree*. Any such edge weighting scheme W induces a metric d_W on X, using the sum of edge weights along paths between pairs of taxa. As is well known, however, a metric d on X need not arise from such a weighting. If $d = d_W$ for some W on T or T^r, then we say d is a *tree metric* on T or T^r with weighting W.

3. Split and clade metrizations of trees

For completeness and perspective on what is to follow, we present two topological metrizations of trees that have been used in other works.

Given an unrooted topological tree T on X, we may assign weights $w(e) = 1$ to all edges e. Although the resulting tree metric on X is just the usual graph-theoretic distance along T, and is essentially that used in [LY11] (see [ADR17]), it may also be interpreted in terms of splits. For this reason, we denote the weighting scheme with all weights 1 by Sp, and say it gives the *split metrization* of T. For $x, y \in X$ the graph-theoretic distance $d_{Sp}(x, y)$ can also be defined as the number of splits $A|B$ displayed on T with $x \in A, y \in B$ or $x \in B, y \in A$. This is a direct consequence of the correspondence between displayed splits and edges on the tree. More informally, $d_{Sp}(x, y)$ is the number of splits displayed on T that *separate* x and y.

For a rooted topological tree T^r on X, assign numbers to the internal nodes of the tree as follows: To the root assign N, to its children that are internal nodes assign $N-1$, to their children that are internal nodes assign $N-2$, and so on, decreasing by 1 for each parent-to-child step. Number all leaves with 0. Then assign weights $w(e)$ as the positive difference of the numbers on the endpoints of e. Thus all internal edges are weighted 1, but terminal edges are weighted at least 2. All leaf-to-root distances are N, so the tree is ultrametric. This is the standard STAR weighting of [LYPE09, ADR13], though here we denote this weighting scheme by Cl, say it gives a *clade metrization* of T^r, and denote the induced metric on X by d_{Cl}. The name is justified by the observation in [ADR13] that for $x, y \in X$,

$$d_{Cl}(x, y) = 2(1 + N - |C_{x,y}|)$$

where $C_{x,y}$ is the set of clades displayed on T^r that contain both x and y. This follows directly by the correspondence between nodes on the path between the tree root and $\mathrm{MRCA}(x, y)$ to clades containing x and y.

Remark 3.1. There are generalizations of the clade metrization, defined in [ADR13], which allow for the weight of an edge to depend on the number of edges between it and the root. These can be used for consistent species tree inference, in generalizations of STAR.

Remark 3.2. One might propose an alternative clade metrization, Cl' defined by assigning unit weights to all edges in a rooted tree. Then the intertaxon distance $d_{Cl'}(x, y)$ is the number of clades displayed on the tree that contain one of x, y but not both. However, $d_{Cl'}$ lacks the ultrametricity of d_{Cl}. It can also be viewed as the split distance on the unrooted tree obtained by attaching a single edge with an extra taxon to the root.

4. Rooted triple metrization of a rooted tree

With T^r a rooted binary phylogenetic tree on X, we may assign edge weights to T^r as follows: First number each node of the tree, including leaves, with the number of taxa descended from it. Leaves are numbered 1, as they are considered their own descendants, and the root is numbered N, the total number of taxa. Then assign weights $w(e)$ as the positive difference of the numbers on the endpoints of e. That is, for any edge $e = (u, v)$ directed away from the root with, u the parent of v, the edge weight is

$$w(e) = |\operatorname{descendants}(u)| - |\operatorname{descendants}(v)|,$$

the decrease in number of descendants across e. We refer to this as the *rooted triple metrization*, due to Theorem 4.3 below, and denote the weighting scheme RT. It results in an ultrametric tree, with the root at distance $N - 1$ from every leaf.

Example 4.1. With the rooted triple metrization, a rooted caterpillar tree

$$(\dots(((a_1, a_2), a_3), a_4), \dots, a_N)$$

will have all internal edges of weight 1. The pendant edges, listed from the cherry toward the root, will have weights $1, 1, 2, 3, 4, \dots, N - 1$.

Example 4.2. With the rooted triple metrization, a balanced tree

$$(\dots((a_1, a_2), (a_3, a_4)), \dots, ((a_{N-3}, a_{N-2}), (a_{N-1}, a_N))\dots)$$

on $N = 2^k$ taxa will have pendant edges of weight 1, and as one moves toward the root, internal edges of weight $2, 4, 8, \dots, 2^{k-1}$.

Theorem 4.3. *Suppose a rooted binary phylogenetic tree T^r is given the rooted triple metrization. Then the resulting tree metric d_{RT} satisfies, for all $x, y \in X$, $x \neq y$,*

$$d_{RT}(x, y) = 2|R_{x,y}| + 2,$$

where $R_{x,y}$ is the set of rooted triples displayed on T of the form $xz|y$ or $yz|x$.

More informally, the distance $d_{RT}(x, y)$ is, up to a simple transformation, the number of rooted triples displayed on T in which x, y are *separated*.

Proof. Let $v = \mathrm{MRCA}(x, y)$, and k be the number of leaf descendants of v (i.e., k is the size of the smallest displayed clade containing x, y. Since $d_{RT}(x, y)$ is the sum of edge weights on the path between x and y, we find that

$$d_{RT}(x, y) = d_{RT}(x, v) + d_{RT}(v, y) = (k - 1) + (k - 1) = 2k - 2.$$

The number of rooted triples of the form $xz|y$ or $yz|x$ is the number of taxa z descended from v, excluding x and y. Thus

$$|R_{x,y}| = k - 2.$$

Eliminating k from these two equations yields the claim. □

The rooted triple metrization can be applied to non-binary trees as well, with edge weights still given by the drop in number of descendants from parent to child nodes. In this setting one one must consider degenerate rooted triples xyz when the 3-taxon tree induced on x, y, and z is a star (unresolved) tree. Minor modifications to the proof above yield the following.

Theorem 4.4. *Suppose a rooted phylogenetic tree T^r is given the rooted triple metrization. Then the resulting tree metric d_{RT} satisfies, for all $x, y \in X$, $x \neq y$,*

$$d_{RT}(x, y) = 2|\tilde{R}_{x,y}| + 2, \tag{1}$$

where $\tilde{R}_{x,y}$ is the set of rooted triples displayed on T of the form $xz|y$, $yz|x$, or xyz.

Remark 4.5. Combining Theorems 4.3 or 4.4 with the fact that one can determine a metric tree from its intertaxon distances gives an alternative proof of the well-known result that the collection of rooted triples displayed on a tree determines the rooted tree.

5. Quartet metrization of a unrooted tree

Let T be an unrooted binary tree on taxa X. Each internal edge of T determines a partition of X into 4 non-empty blocks, X_1, X_2, X_3, X_4 where the split associated to the edge is $X_1 \cup X_2 | X_3 \cup X_4$, and the splits associated to the 4 adjacent edges all have an X_i as one split set. We will refer to this partition as the *quartet partition* associated to an internal edge, and denote it by $X_1, X_2|X_3, X_4$. Assign an internal edge e with quartet partition $X_1, X_2|X_3, X_4$ the weight

$$w_Q(e) = |X_1||X_2| + |X_3||X_4|.$$

For a pendant edge to leaf x, the non-leaf endpoint determines a tripartition of the taxa, $\{x\}, X_1, X_2$. Assign to such an edge e the weight

$$w_Q(e) = |X_1||X_2|.$$

We refer to this as the *quartet metrization*, due to Theorem 5.3 below, and denote the weighting scheme Q. Trees with the quartet metrization are usually not ultrametric, as examples show.

Example 5.1. A unrooted caterpillar tree

$$(\dots(((a_1, a_2), a_3), a_4), \dots, a_N)$$

will have internal edges inducing quartet partitions $X_1, X_2|X_3, X_4$ with sets of size

$$|X_1| = k - 1,\ |X_2| = 1,\ |X_3| = 1,\ |X_4| = N - k - 1,$$

for $k = 2, 3, \dots N - 2$. Under the quartet metrization, the internal edge weights will thus all be $N - 2$. The pendant edges to taxa a_1, a_2, a_{N-1}, a_N, will also have weights $N - 2$.

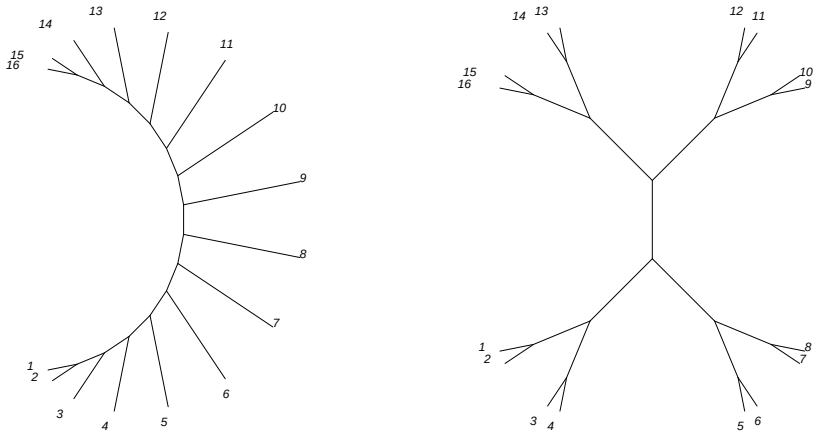

FIGURE 5.1. 16-taxon caterpillar and balanced trees, with edge lengths given by the quartet metrization

Pendant edges to taxa a_k, for $k = 3, \dots N-2$, will have weights $(k-1)(N-k)$. A 16-taxon illustration is shown in Figure 5.1.

Example 5.2. An unrooted balanced tree

$$(\dots((a_1, a_2), (a_3, a_4)), \dots, ((a_{N-3}, a_{N-2}), (a_{N-1}, a_N))\dots)$$

on $N = 2^k$ taxa will have pendant edges of weight $N-2$. With $\ell \in 1, 2, \dots, k-1$ denoting the minimal number of edges needed to connect a given internal edge to a leaf, the central internal edge, for which $\ell = k-1$, has weight $2^{k-2} \cdot 2^{k-2} + 2^{k-2} \cdot 2^{k-2} = 2^{2k-3}$, while other internal edges, with $1 \leq \ell \leq k-2$, are of weight

$$(2^{\ell-1})(2^{\ell-1}) + (2^{\ell})(2^k - 2^{\ell+1}) = 2^{k+\ell} - 7 \cdot 2^{2\ell-2}.$$

A 16-taxon illustration is shown in Figure 5.1.

Theorem 5.3. *Suppose an unrooted binary phylogenetic tree T has been given the quartet metrization. Then the resulting tree metric d_Q satisfies for all $x, y \in X$, $x \neq y$,*

$$d_Q(x, y) = 2|Q_{x,y}| + 2N - 4 \tag{2}$$

where $Q_{x,y}$ is the set of quartets displayed on T of the form $xz|yw$.

More informally, the distance $d_Q(x, y)$ is, up to a simple transformation, the number of quartets displayed on T in which x, y are *separated*.

Although Theorem 5.3 can be deduced from Theorem 4.3 by summing its formula over all placements of the root on pendant edges of T, a more direct argument for it will be given here.

Proof. Fix taxa $x \neq y$, and let P denote the path in T between them. Any node $u \neq x, y$ on P determines a tripartition of the taxa $X = A_u \sqcup B_u \sqcup C_u$ according to the connected components resulting from the deletion of u and its incident edges from the tree. We may require $x \in A_u$ and $y \in B_u$, so the set C_u contains all those taxa w for which a path from w to x or y joins P at u.

Now any quartet $xw|yz$ separating x and y that is displayed on the tree T determines a node v on P by where the path from w to y joins P. Furthermore, the number of quartets separating x and y that are displayed on T and determine the same node v in this way is $|C_v|(|B_v| - 1)$, so

$$|Q(x,y)| = \sum_{\substack{v \text{ on } P \\ v \neq x,y}} (|C_v||B_v| - |C_v|) = \sum_{\substack{v \text{ on } P \\ v \neq x,y}} |C_v||B_v| - N + 2. \tag{3}$$

Interchanging the roles of x and y we also find

$$|Q(x,y)| = \sum_{\substack{v \text{ on } P \\ v \neq x,y}} |C_v||A_v| - N + 2, \tag{4}$$

and adding equations(3) and (4) yields

$$2|Q(x,y)| = \sum_{\substack{v \text{ on } P \\ v \neq x,y}} |C_v||A_v| + \sum_{\substack{w \text{ on } P \\ w \neq x,y}} |C_w||A_w| - 2N + 4.$$

Writing the right hand side as a sum over edges $e = (v, w)$ in P shows

$$2|Q(x,y)| = \sum_{e \text{ on } P} w_Q(e) - 2N + 4 = d_Q(x,y) - 2N + 4,$$

which yields the claim. □

Remark 5.4. Proposition 9 of [BD86] shows that $|Q(x,y)| + 1$ for $x \neq y$ yields a tree metric on T, which is equivalent to the right hand side of equation (2) defining a tree metric on T. However, edge weights associated to the tree metric are not investigated in that paper. Moreover, applications of the result to tree inference, such as those discussed in Sections 6 and 7 of this work, seem not to have been pursued in intervening years.

While the quartet metrization is defined above only for binary trees, it can be extended to non-binary trees as follows: Suppose an edge $e = \{u, v\}$ is incident to m_u other edges e_i at u, and m_v other edges $\tilde{e}_j$ at v. Note that for a leaf v we allow $m_v = 0$. Now let $X_1, X_2, \ldots X_{m_u}, \tilde{X}_1, \tilde{X}_2, \ldots, \tilde{X}_{m_v}$ be the partition of X with X_i (respectively $\tilde{X}_j$) the set of taxa connected to e by a path through e_i (respectively $\tilde{e}_j$). Then assign to e the edge weight

$$w_Q(e) = \sum_{1 \leq i < i' \leq m_u} |X_i||X_{i'}| + \sum_{1 \leq j < j' \leq m_u} |\tilde{X}_j||\tilde{X}_{j'}|.$$

Interpreting an empty sum as 0, this agrees with the definition above for binary trees.

If one allows for degenerate quartets $xyzw$ to be displayed on a tree when the taxa x, y, z, w induce a star (unresolved) tree, this leads to the following generalization of Theorem 5.3, with a similar proof.

Theorem 5.5. *Suppose an unrooted phylogenetic tree T has been given the quartet metrization. Then the resulting tree metric d_Q satisfies for all $x, y \in X$, $x \neq y$,*

$$d_Q(x, y) = 2|\tilde{Q}_{x,y}| + 2N - 4$$

where $\tilde{Q}_{x,y}$ is the set of quartets displayed on T of the form $xz|yw$ or $xyzw$.

Remark 5.6. Combining Theorems 5.3 or 5.5 with the fact that one can determine a metric tree from its intertaxon distances gives an alternative proof of the well-known result that the collection of quartets displayed on a tree determines the unrooted tree.

Remark 5.7. As a heuristic, ASTRAL-II uses a similarity on taxa that counts quartets *not* separating two taxa on a gene tree [MW15]. By Theorem 5.3 this is essentially equivalent to the quartet metrization for that gene tree. While ASTRAL-II's goal is species tree inference, the details of how it uses this similarity are, however, quite different from what is presented in Section 7 below for statistically consistent inference of a species tree.

6. Quartet Distance Supertree

Theorems 4.3 and 5.3 lead to new supertree methods for finding a tree from certain collections of rooted triples or quartets. We present this fully for quartets, indicating the small modifications for rooted triples in a remark.

Suppose we are given a collection $\mathcal{Q}$ of unweighted quartets on a set of taxa X. We take the viewpoint that most of the given quartets show the correct phylogenetic relationship between the taxa, though some are in error. Ideally $\mathcal{Q}$ contains exactly one quartet for each subset of 4 taxa.

We choose a distance-based method M of tree building or selection that when applied to a tree metric on an unrooted tree T returns T, even if T is not ultrametric. We further require that its output topology is robust to small errors in the input distances at a tree metric. Possible choice for M include NJ and BioNJ (but not UPGMA) for tree building [SN87, Gas97], and Balanced Minimum Evolution for tree selection. In practice, a heuristic implementation known to perform well, such as FastME [DG02], may be chosen.

Algorithm 6.1. (QDS/M) Quartet Distance Supertree with distance method M
Input: A collection $\mathcal{Q}$ of quartets on taxa in X

(1) For each pair $x, y \in X$ of taxa, $x \neq y$, count the number $q(x, y)$ of quartets in $\mathcal{Q}$ separating x, y, and define the distance $\hat{d}_Q(x, y) = 2q(x, y) + 2N - 4$.
(2) Use the distance method M to build or select an unrooted tree from $\hat{d}_Q$.

Remark 6.2. If $\mathcal{Q}$ omits quartets for a small number of 4-taxon subsets, or includes a small number of redundant ones, one could view this an additional error, and modify the algorithm slightly. For instance, one might use all quartets on a given set of 4 taxa by weighting them by relative frequency in the counting. Omitted 4-taxon subsets might be left out of counting when determining intertaxon distances, or treated as all 3 possible quartets on those taxa, each weighted by 1/3, in counting. However, we have no theoretical

justification for these suggestions, and note they conflict with how this algorithm should be used for species tree inference as discussed in the next section.

Remark 6.3. For Rooted Triple Distance Supertree with M, one instead counts the number $r(x, y)$ of rooted triples in a set $\mathcal{R}$ that separate x, y, and defines $\hat{d}_{RT} = 2r(x, y) + 2$. The method M can now be chosen to assume ultrametricity (e.g., UPGMA), since $\hat{d}_{RT}$ approximates the ultrametric tree metric d_{RT}. If so, then a rooted tree will be returned.

We refer to the method of Algorithm 6.1 as *Quartet Distance Supertree* (QDS), and the variant for Rooted Triples as *Rooted Triple Distance Supertree* (RTDS), but emphasize that for a complete specification it is necessary to also indicate the distance method M used for tree construction or selection. If none of the quartets in $\mathcal{Q}$ are erroneous or omitted, these algorithms recover the correct tree. However, how much error, and of what form, can occur in $\mathcal{Q}$ with the desired tree still being accurately recovered may depend on the particular distance method M used. Since theoretical guarantees on toleration of error by distance methods tend to be much weaker than results seen in simulation studies, performance of QDS/M will need to be judged through simulation.

The following simulations, performed in R using the ape package [PCS04], give a first indiction of the performance of QDS. Using the two extreme topologies of caterpillar and balanced trees on 16 taxa, the set of all displayed quartets was formed. Error was then introduced into the quartets in one of two ways. In the first scenario, for choices of probability $0 < p \leq .5$ of quartet error, each quartet was modified with probability p to one of the two alternatives on the same taxa (with equal probability). In the second scenario, for choices of probability $0 < p \leq .9$, the quartet was removed from the set. For each of these modified quartet sets, QDS/NJ was used to construct a tree. In the second scenario, omitted quartets were simply left out of the counting that determines intertaxon distances. The Robinson-Foulds distance was then computed between the inferred QDS/NJ tree and the original tree. This was repeated 100 times, with results summarized in the plots of Figure 6.1. Similar results (not shown) were obtained using the FastME heuristic for balanced minimum evolution in place of NJ.

These plots show that even with about a quarter of the quartets incorrect, on average the correct tree was recovered to within an RF distance of 2 (i.e, all but 1 of the 13 non-trivial splits were recovered correctly) for the caterpillar tree. And even with about half of the caterpillar's quartets omitted, results were similarly accurate. The balanced tree topology was even more robustly recovered than the caterpillar tree, allowing quite large amounts of quartet error under both scenarios.

Of course one should interpret these results cautiously, as empirical quartet error may not have the simple form simulated here. In empirical settings it is unlikely that all quartets would be equally likely to be incorrect or omitted, or that in the case of an incorrect quartet that both alternatives would be equally likely. Nonetheless, these simulations strongly indicate the need for more realistic simulation studies to investigate performance.

Remark 6.4. An obvious drawback of QDS for general tree inference from quartets is that its theoretical basis assumes one has a quartet in $\mathcal{Q}$ for every subset of 4 taxa, and no weights can be supplied expressing relative confidence in those quartets. This differs

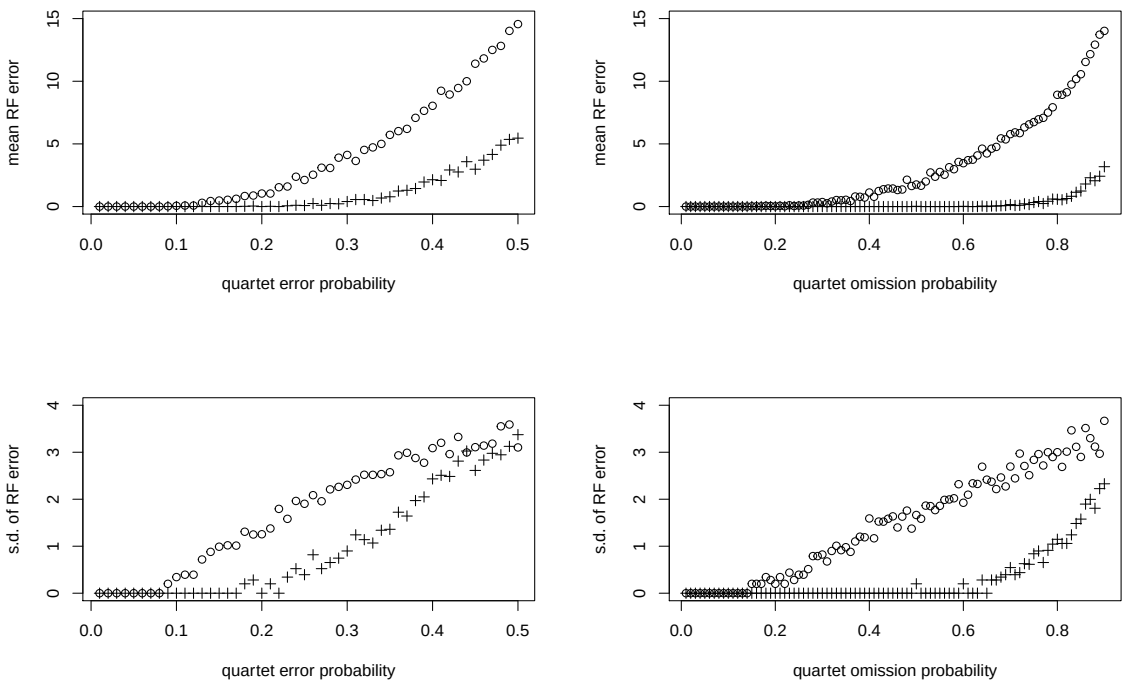

FIGURE 6.1. Performance of QDS/NJ in simulation under under two scenarios, erroneous quartets (left column) and omitted quartets (right column), as described in the text. Circles (∘) denote the 16-taxon caterpillar tree and pluses (+) the 16-taxon balanced tree. The horizontal axes on the left plots gives the probability that a true quartet is replaced with an alternative on the same taxa in forming the quartet set. On the right plots, it gives the probability a quartet is omitted from the quartet set. The vertical axes are the mean (top row) and standard deviation (bottom row) of the RF distance between the original tree and the inferred tree over 100 trials.

from the maximum quartet consistency framework in which one seeks to maximize an objective function expressing the total weights of quartets displayed on the tree. While the above simulations suggest uniformly missing quartets may be of less concern, confidence weighting seems to be desirable, at least with quartets inferred by Maximum Likelihood, [RG01]. However, in some applications, and especially for species tree inference from gene trees as described in the next section, these aspects of QDS may not be a great disadvantage.

Remark 6.5. It is possible that new distance methods could be developed that are more finally tuned to QDS than those existing now. Since the distance $\hat{d}_Q$ approximates distances on an unknown tree T endowed with the quartet metrization, a tree building or

selection method that takes that specific metrization into account may improve performance. Current distance methods are general, making no assumption about a trees edge lengths as related to its topology.

7. Species Tree inference by Quartet Distance Consensus

We next show how QDS and RTDS can be applied to the problem of inferring a species tree from a collection of gene trees. This provides new consensus methods that are statistically consistent under the multispecies coalescent model [DDBR09]. For simplicity, we focus on the application of QDS.

For inference from multilocus sequence data, this can be used in a two-step procedure in which gene trees are first inferred from gene sequences, and then these inferred gene trees are treated as data for inference of a species tree. As is common for such two-stage schemes, the second stage of this method is provably statistically consistent, in the sense that if the gene trees were sampled without error under the multispecies coalescent model, then as the number of gene trees increases the probability of inferring the correct species tree approaches 1. In practice, however, there may be some inference error in the gene trees, as well as violations of the coalescent model, such as horizontal gene transfer.

For the algorithm, we assume we already have in hand a collection of binary gene trees on X, but allow some missing taxa on each tree. However, for good performance it is desirable that each 4-taxon subset appears on many of the gene trees.

Algorithm 7.1. (QDC/M) Quartet Distance Consensus with distance method M
Input: A collection of gene trees on taxa X

(1) For each subset of four taxa $x, y, z, w \in X$, determine the *dominant* (i.e., most frequent) quartet $xy|zw$, $xz|yw$, or $xw|yz$ displayed on the gene trees. In the case of a tie, choose from the most frequent uniformly at random.
(2) With $\mathcal{Q}$ the set of dominant quartets, apply QDS with M.

Modifications to formulate Rooted Triple Distance Consensus (RTDC) are straightforward.

While QDC is similar to the clade-distance based STAR [LYPE09, ADR13] and split-distance based NJ_{st} (a.k.a. ASTRID and U-STAR) [LY11, VW15, ADR17], both of those average distances across gene trees, while QDC/M instead chooses the dominant quartets across gene trees to define a distance. Note that Rooted Triple Consensus [EESvH08] and BUCKy [LKDA10] similarly chose the dominant rooted triple or quartet for species tree inference, but then use different algorithms for finding a larger tree.

Indeed, under the multispecies coalescent model on a binary species tree the dominant quartet across gene trees is known to agree with the species tree quartet with probability approaching 1 as the sample size (number of independent loci) goes to infinity [ADR11]. This fact ensures statistically consistent inference of the species tree when the algorithm is applied to a sample of true gene trees without missing taxa. Even if gene trees have missing taxa, the algorithm retains statistical consistency under an assumption that for each choice of four elements of X, as the number of loci goes to infinity, the number of gene trees on which the four appears also goes to infinity.

The ability to deal with missing taxa is potentially an advantage of species tree inference by QDC over the NJ$_{st}$/ASTRID/U-STAR approach. Even though simulations have shown good performance of ASTRID with taxa missing from gene trees [VW15], there is as yet no theoretical argument supporting consistency in that circumstance, and it is unclear whether the patterns of missing taxa that one might have in an empirical study are captured by the simulations. On the other hand, ASTRID is likely to be faster than any implementation of QDC, since QDC appears to require consideration of each of the $\binom{n}{4}$ subsets of taxa on each tree to determine dominant quartets, while ASTRID can simply average distance matrices.

For a first look at the possible performance of QDC for species tree inference, it was applied to simulated data sets from [BMW14], which were also analyzed in [VW15]. These data sets contain 20 replicate samples of 200 multispecies coalsecent gene trees, for each of two 11-taxon species trees, one corresponding to a high level of ILS, and one a very high level, in the terminology of [VW15]. In addition, for each replicate the data sets contain 200 estimated gene trees, obtained by maximum likelihood inference from sequences simulated under a standard substitution model on each sampled gene tree. Roughly following the analysis of Figure 9 of [VW15], both ASTRID and QDC were used to infer species trees from 10, 25, 50, 100, or 200 sampled or estimated gene trees for each of the 20 replicates. Mean Robinson-Foulds distances from the true species tree for the 20 replicates were computed, and plotted in Figure 7.1 These plots suggest that inference with QDC may be slightly more accurate than with ASTRID. Of particular interest, though, at the high level of ILS QDC appears to be more robust to gene tree inference error than is ASTRID. Of course this glimpse at QDC's performance primarily indicates much more extensive investigation is needed.

Remark 7.2. The algorithm above implicitly assumes the species tree is binary, so that for every choice of four taxa there will be a single most probable quartet. For non-binary species trees, one may instead have all three quartets equiprobable (but not have a two-way tie for most probable) in the limit as sample size goes to infinity. Using Theorem 5.5 one could modify the algorithm to allow for non-binary species trees, making some choice of cutoff for judging "near equality" in quartet frequencies.

Remark 7.3. We further note that averaging the quartet distances, or rooted triple distances, across gene trees as is done in STAR and U-STAR would not lead to consistency under the coalescent model. In fact, an example is already given in [ADR13] in which an inconsequential variant of the rooted triple metrization averaged across gene trees is shown to exactly fit an incorrect species tree for infinite sample size. A similar example for the quartet metrization is as follows: Consider the rooted caterpillar species tree $(((((a,b){:}x,c){:}y,d){:}z,e){:}w,f)$ with $x,z,w=\infty$, $y=0$. Then under the multispecies coalescent model the gene trees $(((((a,b),c),d),e),f)$, $(((((a,b),d),c),e),f)$, and $((((a,b),(c,d)),e),f)$ each have probability 1/3, and all others have probability 0. Unrooting these gene trees and applying the quartet metrization, with alphabetical ordering

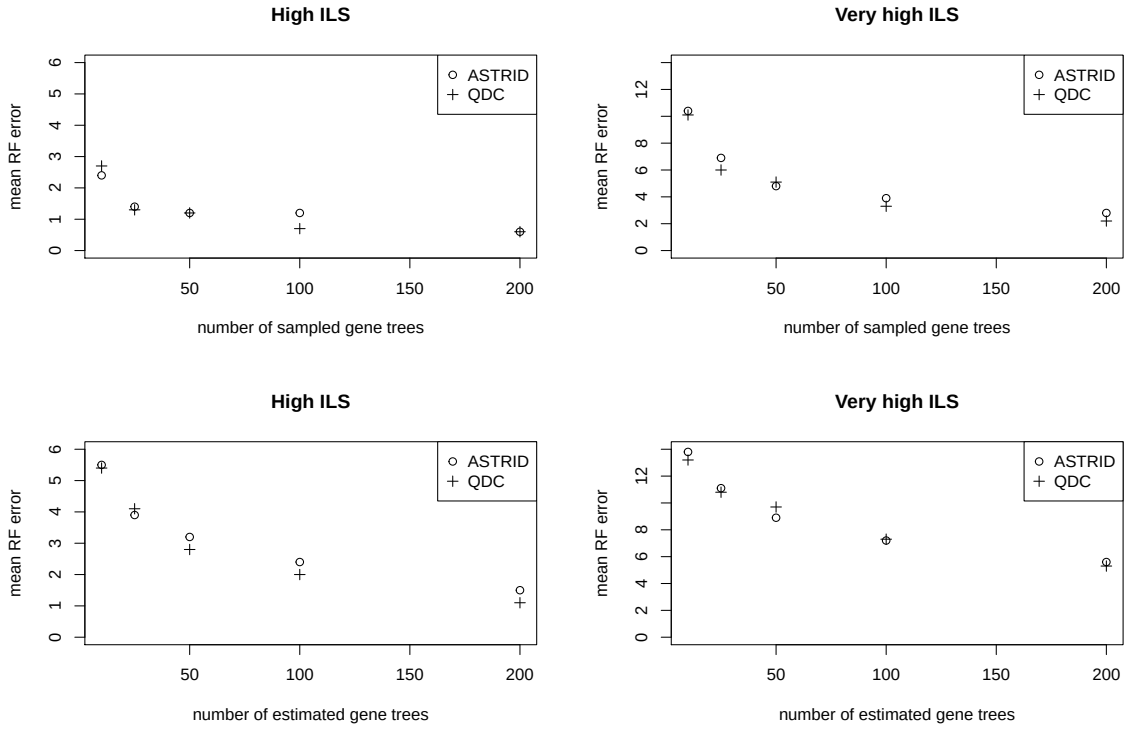

FIGURE 7.1. Performance of ASTRID and QDC/FastME on simulated data of [BMW14] under under two scenarios of high and very high levels of ILS on a species tree with 10 taxa and an outgroup. For the top plots, input to both algorithms were gene trees sampled from the coalescent, with no addiitional error. For the bottom plots, input was gene trees inferred by maximum likelihood from sequences simulated on the sampled gene trees. Note differing vertical scales.

of taxa we obtain the three distance matrices

$$\begin{pmatrix} 0 & 8 & 14 & 18 & 20 & 20 \\ & 0 & 14 & 18 & 20 & 20 \\ & & 0 & 16 & 18 & 18 \\ & & & 0 & 14 & 14 \\ & & & & 0 & 8 \\ & & & & & 0 \end{pmatrix}, \begin{pmatrix} 0 & 8 & 18 & 14 & 20 & 20 \\ & 0 & 18 & 14 & 20 & 20 \\ & & 0 & 16 & 14 & 14 \\ & & & 0 & 18 & 18 \\ & & & & 0 & 8 \\ & & & & & 0 \end{pmatrix}, \begin{pmatrix} 0 & 8 & 18 & 18 & 18 & 18 \\ & 0 & 18 & 18 & 18 & 18 \\ & & 0 & 8 & 18 & 18 \\ & & & 0 & 18 & 18 \\ & & & & 0 & 8 \\ & & & & & 0 \end{pmatrix}.$$

Weighting the matrices by 1/3 and summing yields

$$\begin{pmatrix} 0 & 8 & 50/3 & 50/3 & 58/3 & 58/3 \\ & 0 & 50/3 & 50/3 & 58/3 & 58/3 \\ & & 0 & 40/3 & 50/3 & 50/3 \\ & & & 0 & 50/3 & 50/3 \\ & & & & 0 & 8 \\ & & & & & 0 \end{pmatrix},$$

which exactly agrees with distances on the unrooted tree

$$((a{:}4, b{:}4){:}17/3, (c{:}20/3, d{:}20/3){:}1/3, (e{:}4, f{:}4){:}17/3).$$

Note this does not have the same unrooted topology as the species tree.

To construct an example with a binary species tree with no zero or infinite edge lengths, we perturb the above one slightly. If the distances on the original species tree are chosen so x, z, w are very large and y is very small but positive, the average of the gene tree quartet metrizations will change only slightly. Thus it cannot fit the topology of the species tree.

Finally, note that by using only topological information from gene trees, QDC and RTDC, like STAR and U-STAR, base their inference on those features of gene trees which are usually considered to be those most robustly inferred from sequence data. While metric information on gene trees could carry more information, it is also more sensitive to choice of the substitution model used in inferring the gene trees, and requires further modeling assumptions to relate gene tree and species tree time scales.

8. Discussion

To place this work in the context of other quartet supertree methods, note that the most common framework in quartet methods of inferring trees, maximum quartet consistency, has been to minimize an objective function measuring conflict between the given quartets and the tree, through a search over all possible trees. (An alternative tree construction approach is given in [MXZ08].) In practice, heuristic searches are usually necessary to address the optimization problem, and the number of taxa may be limited in order to achieve acceptable performance and runtimes [Bau92, Rag92, SvH96, SR10, SSLW11, ACS15, MW15]. As reasonable as this broad framework is, however, it is important to remember that the objective functions used are not ones deduced from theory. In fact, no such theory is even possible without an explicit model of error in the quartets, and it does not appear any attempt has been made to justify current approaches in such a way. Instead, simulations which incorporate inference error in the quartets are used for evaluation and comparison of methods.

A rather different notion of fitting a tree to quartets underlies QDS/M, whether the distance method M is a tree building algorithm or optimization of a distance-based objective function. Unfortunately no current theory can guide us as to whether this is better or worse than previous approaches. Extensive simulation studies are needed to judge the

practical effectiveness of the new methods proposed here. Moreover, since the quartet error involved in different applications may have different features, simulations studies must reflect this and be targeted at specific applications. For instance, the effectiveness of QDC for species tree inference from full N-taxon gene trees inferred by Maximum Likelihood (ML) may be different from that of inference by QDS of a single gene tree from quartet trees inferred by ML.

Quartet methods have played a role in recent progress in phylogenetics in using algebraic methods for tree inference from sequence data [CK14, CK15, CFS07, FSC16], and the ideas presented here may be useful for those applications. Using these methods one can infer a quartet tree very quickly under very general models. However, technical issues complicate inference of larger trees directly. If QDS works well with the quartet trees these methods produce, then the significant advantages they offer, in speed and the generality of the underlying substitution model, may be broadened to include quick inference of large trees as well.

For the specific problem of inference of species trees from gene trees, rooted triple and quartet approaches has been taken before [EESvH08, LKDA10, MW15]. However, the difficulty of maximizing the number of triple or quartets displayed on an inferred tree required the use of clever heuristics to achieve acceptable run times. Despite this, the quartet-based ASTRAL-II runs considerably slower than ASTRID, while only having comparable accuracy [VW15]. As ASTRID is based on the split metrization of gene trees, its speed arises in part from the ability to use a fast distance method to obtain the tree. A distance method such as FastME is able to conduct an effective heuristic search for an optimal tree (that is, optimal by an appropriate distance criterion) much more rapidly than has been achieved by any non-distance criterion. While QDC can similarly take advantage of fast distance-based optimization, one should expect that well-written software implementations of it will still be slower than ASTRID, due to the computational time for determining dominant quartets. If QDC leads to even small improvements in accuracy over ASTRID, though, that computational time may be justified.

9. Acknowledgments

This work was supported by the National Institutes of Health grant R01 GM117590, awarded under the Joint DMS/NIGMS Initiative to Support Research at the Interface of the Biological and Mathematical Sciences. Thanks to Elizabeth Allman and Mike Steel, for comments and discussion on an earlier draft.

References

[ACS15] E. Anvi, R. Cohen, and S. Snir. Weighted quartets phylogenetics. *Syst. Biol.*, 64(2):233–242, 2015.

[ADR11] E. S. Allman, J. H. Degnan, and J. A. Rhodes. Identifying the rooted species tree from the distribution of unrooted gene trees under the coalescent. *J. Math. Biol.*, 62(6):833–862, 2011.

[ADR13] E. S. Allman, J. H. Degnan, and J. A. Rhodes. Species tree inference by the STAR method, and generalizations. *J. Comput. Biol.*, 20(1):50–61, 2013.

[ADR17] Elizabeth S. Allman, James H. Degnan, and John A. Rhodes. Species tree inference from gene splits by Unrooted STAR methods. *IEEE/ACM Trans. Comput. Biol. Bioinf.*, to appear, 2017.

[Bau92] B. R. Baum. Combining trees as a way of combining data sets for phylogenetic inference, and the desirability of combining gene trees. *Taxon*, 41(3–10), 1992.

[BD86] H.-J. Bandelt and A. Dress. Reconstructing the shape of a tree from observed dissimilarity data. *Adv. Appl. Math.*, 7:309–343, 1986.

[BMW14] M.S. Bayzid, S. Mirarab, and T. Warnow. Weighted statistical binning: enabling statistically consistent genome-scale phylogenetic analyses. *PLOS One*, 2014.

[Bry03] D. Bryant. A classification of consensus methods for phylogenetics. *DIMACS series in discrete mathematics and theoretical computer science*, 61:163–184, 2003.

[CFS07] M. Casanellas and J. Fernández-Sánchez. Performance of a new invariants method on homogeneous and non-homogeneous quartet trees. *Mol. Biol. Evol.*, 24(1):288–293, 2007.

[CK14] J. Chifman and L. Kubatko. Quartet inference from SNP data under the coalescent model. *Bioinformatics*, 30(23):3317–3324, 2014.

[CK15] J. Chifman and L. Kubatko. Identifiability of the unrooted species tree topology under the coalescent model with time-reversible substitution processes, site-specific rate variation, and invariable sites. *J. Theor. Biol.*, 374:35–47, 2015.

[DDBR09] J. H. Degnan, M. DeGiorgio, D. Bryant, and N. A. Rosenberg. Properties of consensus methods for inferring species trees from gene trees. *Syst. Biol.*, 58(1):35–54, 2009.

[DG02] R. Desper and O. Gascuel. Fast and accurate phylogeny reconstruction algorithms based on the minimum-evolution principle. *J. Comput. Biol.*, 9(5):687–705, 2002.

[EESvH08] G.B. Ewing, I. Ebersberger, H.A. Schmidt, and A. von Haeseler. Rooted triple consensus and anomalous gene trees. *BMC Evol. Biol.*, 8:118, 2008.

[FSC16] J. Fernández-Sánchez and M. Casanellas. Invariant versus classical quartet inference when evolution is heterogeneous across sites and lineages. *Syst. Biol.*, 65(2):280–291, 2016.

[Gas97] O. Gascuel. BIONJ: an improved version of the NJ algorithm based on a simple model of sequence data. *Mol. Biol. Evol.*, 14:685–695, 1997.

[LC97] J.F. Lapointe and G. Cucumel. The average consensus procedure: combination of weighted trees containing identical or overlapping sets of taxa. *Syst. Biol.*, 46(2):306–312, 1997.

[LKDA10] B. R. Larget, S. K. Kotha, C. N. Dewey, and C. Ané. BUCKy: Gene tree / species tree reconciliation with Bayesian concordance analysis. *Bioinformatics*, 26:2910–11, 2010.

[LY11] L. Liu and L. Yu. Estimating species trees from unrooted gene trees. *Syst. Biol.*, 60:661–667, 2011.

[LYPE09] L. Liu, L. Yu, D. K. Pearl, and S. V. Edwards. Estimating species phylogenies using coalescence times among sequences. *Syst. Biol.*, 58:468–477, 2009.

[MW15] Siavash Mirarab and Tandy Warnow. ASTRAL-II: coalescent-based species tree estimation with many hundreds of taxa and thousands of genes. *Bioinformatics*, 31:i44–i52, 2015.

[MXZ08] B. Ma, L. Xin, and K. Zhang. A new quartet approach for reconstructing phylogenetic trees: quartet joining method. *J. Comb. Optim.*, 16(3):293–306, 2008.

[PCS04] E. Paradis, J. Claude, and K. Strimmer. APE: analyses of phylogenetics and evolution in R language. *Bioinformatics*, 20:289–290, 2004.

[Rag92] M. A. Ragan. Phylogenetic inference based on matrix representation of trees. *Mol. Phylo. Evol.*, 1(53–58), 1992.

[RG01] V. Ranwez and O. Gascuel. Quartet-based phylogenetic inference: Improvements and limits. *Mol. Biol. Evol.*, 18(6):1103–1116, 2001.

[SN87] N. Saitou and M. Nei. The neighbor-joining method: A new method for reconstruction of phylogenetic trees. *Mol. Biol. Evol.*, 4:406–425, 1987.

[SR10] S. Snir and S. Rao. Quartets MaxCut: a divide and conquer quartets algorithm. *IEEE/ACM Trans. Comput. Biol. Bioinf.*, 7(4):704–718, 2010.

[SSLW11] M. S. Swenson, R. Suri, C. R. Linder, and T. Warnow. An experimental study of quartets MaxCut and other supertree methods. *Alg. Mol. Biol.*, 6(1):7, 2011.

[SvH96] K. Strimmer and A. von Haeseler. Quartet puzzling: A quartet maximum-likelihood method for reconstructing tree topologies. *Mol. Biol. Evol.*, 13(7):964–969, 1996.

[VW15] P. Vachaspati and T. Warnow. ASTRID: Accurate Species TRees from Internode Distances. *BMC Genomics*, 16(Suppl 10):S3, 2015.

Department of Mathematics and Statistics, University of Alaska Fairbanks, 99775
E-mail address: j.rhodes@alaska.edu